



NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

Modeling and Analysis of Uncertain Time-Critical Tasking Problems (UTCTP)

by

Donald P. Gaver, Jr.
Patricia A. Jacobs
Gennady Samorodnitsky

November 2003

Approved for public release; distribution is unlimited.

Prepared for: Office of Naval Research, Math/Computer/Information Sciences
800 North Quincy Street, Ballston Tower 1, Code 311
Arlington, VA 22217-5660

**NAVAL POSTGRADUATE SCHOOL
MONTEREY, CA 93943-5000**

RADM David R. Ellison
Superintendent

Richard Elster
Provost

This report was prepared for and funded by Office of Naval Research,
Math/Computer/Information Sciences, 800 North Quincy Street, Ballston Tower 1,
Code 311, Arlington, VA 22217-5660

Reproduction of all or part of this report is authorized.

This report was prepared by:

DONALD P. GAVAR, JR.
Distinguished Professor of
Operations Research

PATRICIA A. JACOBS
Professor of Operations Research

Reviewed by:

GENNADY SAMORODNITSKY
Professor of Operations Research

LYN R. WHITAKER
Associate Chairman for Research
Department of Operations Research

Released by:

JAMES N. EAGLE
Chairman
Department of Operations Research

LEONARD A. FERRARI, Ph.D.
Associate Provost and Dean of Research

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE November 2003	3. REPORT TYPE AND DATES COVERED Technical Report	
4. TITLE AND SUBTITLE: Modeling and Analysis of Uncertain Time-Critical Tasking Problems (UTCTP)			5. FUNDING NUMBERS	
6. AUTHOR(S) Donald P. Gaver, Jr., Patricia A. Jacobs, Gennady Samorodnitsky				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			8. PERFORMING ORGANIZATION REPORT NUMBER NPS-OR-03-005	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Office of Naval Research, Math/Computer/Information Sciences 800 North Quincy Street, Ballston Tower 1, Code 311 Arlington, VA 22217-5660			10. SPONSORING / MONITORING AGENCY REPORT NUMBER N/A	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) <p>We consider modeling and operational analysis of a generic asymmetric service-system situation in which (a) Red agents, potentially threatening, but possibly requiring assistance, arrive according to some partially known and possibly changing pattern in time and space; and (b) Reds are <i>impatient</i>: have effectively limited unknown deadlines or times of availability for Blue service, i.e., detection, classification, and attack in a military setting, or emergency assistance in other settings, such as medical care.</p> <p>We discuss various service options by Blue service agents and devise several approximations allowing one to compute efficiently proportions of tasks of different types that are successfully served or, more generally, if different rewards are associated with different types of tasks, the percentage of the possible reward gained.</p> <p>We suggest a heuristic policy for a Blue server to select the next task to perform and to decide how much time to allocate to that service. We discuss this and related policies for a number of specific examples.</p>				
14. SUBJECT TERMS Queueing, Random Task Deadlines, Time Critical Tasking, Myopic scheduling policy			15. NUMBER OF PAGES 36	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

Modeling and Analysis of Uncertain Time-Critical Tasking Problems (UTCTP)

Donald P. Gaver

Patricia A. Jacobs

Operations Research Department,
Naval Postgraduate School, Monterey, CA 93943

Gennady Samorodnitsky

School of Operations Research and Industrial Engineering,
Cornell University, Ithaca, NY 14853

ABSTRACT

We consider modeling and operational analysis of a generic asymmetric service-system situation in which (a) Red agents, potentially threatening, but possibly requiring assistance, arrive according to some partially known and possibly changing pattern in time and space; and (b) Reds are *impatient*: have effectively limited unknown deadlines or times of availability for Blue service, i.e., detection, classification, and attack in a military setting, or emergency assistance in other settings, such as medical care.

We discuss various service options by Blue service agents and devise several approximations allowing one to compute efficiently proportions of tasks of different types that are successfully served or, more generally, if different rewards are associated with different types of tasks, the percentage of the possible reward gained.

We suggest a heuristic policy for a Blue server to select the next task to perform and to decide how much time to allocate to that service. We discuss this and related policies for a number of specific examples.

1. *The Problem*

Consider modeling and operational analysis of a generic asymmetrical *service-system situation* in which (a) Red *agents*, such as, military or facility-destructive

hostile threats, arrive according to some partially known and possibly changing pattern in time and space; and (b) Reds have effectively limited unknown deadlines, or times of availability for Blue *service*, i.e., detection, classification, and attack in a military setting. Cases of known deadlines are important and somewhat analogous; see Lehoczky (1996, 1997a, 1997b) and Doytchinov *et al* (2001).

Think of the Reds as presenting tasks to be performed, or to be subjects of service. In a military context Reds may be perceived as enemy targets, but in a medical emergency room setting they're arriving casualties. In a call center they are requests for information; see Becker *et al* (2000). In a Homeland Security (HLS) scenario a Red may be a container ship approaching a port, or a truck approaching a border, either possibly carrying explosives or chemical-biological offensive agents. We consider the Blue problem of processing such Red tasks effectively and efficiently under time constraints and limited information, hence the necessity to control the amount of service given.

Appropriate service effort typically differs between task types; it may not always be completely provided, and may be partial and incomplete, owing to deficiency of time, information or resources. In general, task service is by a Blue force of task-server agents also of various types, possibly varying in number and organization, and at different locations, but which attempt to share information and the service burden. Such complex agent systems are considered elsewhere, using insights provided in this report.

Some General Questions: How to match the Red tasks to currently appropriate Blue servers? How many, and what types of, Blue servers are needed to cope with the range of Red demands? How adequate is Blue service of Red, where "service" here means neutralization of threat (or recovery of endangered isolated personnel, such as a

downed pilot or downed aircraft human and other valuable and sensitive cargo), or stabilization of an injury, or identification and proper “decontamination” of a platform carrying dangerous cargo? There are many examples of the generic and simplified situation we consider.

The Blue military objective is to successfully service as many tasks as possible rather than to minimize queues, while hostile Reds attempt to avoid “service,” at least until they can accomplish their purpose, often to damage Blue. The models presented and analyzed suggest Blue force requirements and capability combinations for confronting specified challenges with acceptable success rates.

To summarize, such service system issues arise ubiquitously in military operations of all kinds, as well as in Homeland Security and in military force protection, emergency management situations, and many natural hazard response scenarios, such as after earthquakes or tidal waves. They also occur in call center design and operation, wherein specialized operators are made available to assist users of new software issues (see Becker *et al* (2000)).

The plan of this paper is as follows. In Section 2, we describe a simple version of the basic problem: A single-type Red task arrival stream confronts a single Blue server that can process just one Red task at a time; success probability is related to allocated processing time, but is considered fixed/constant for a given selected processing mode. The Red tasks each have randomly limited availability time for processing, so if that time for a task exceeds any waiting plus service allocation time then service is delivered with a possibly successful outcome. Otherwise, the task is lost (leaks through defenses, or dies while awaiting treatment, etc.).

In Section 3, the investigation is broadened to include several task types. We suggest approximations to the proportion of tasks that are successfully served. We examine Blue defense's decision options so as to achieve maximum success rate, i.e., minimum leakage probability. These are (a) to select for next service the waiting task with greatest chance of survival to be serviced, and (b) to allocate service resources so as to balance time spent on the currently served task against losing tasks waiting. Control Policies are proposed that may then be refined and evaluated by use of heuristic search procedures such as Genetic Algorithms, and by adaptations of Dynamic Programming.

2. One Red Task-Type vs. Single Blue Service Agent

Consider the simplest case of a single Blue service-providing agent (BSA) confronting a random stream of identical loss-susceptible Red service-requiring agents (RSAs).

Model 1

Assume first that members of the RSA stream arrive at (enter the sector of) a single BSA. Rate of approach is λ ; after arrival the n th-to-arrive RSA has a loss time L_n : unless served within time L_n , the task vanishes. Optionally, $\{L_n\}$ is a sequence of independently identical distributed random variables, with expectation $E[L_n] = \theta^{-1}$, but more generality is possible. For convenience, L_n may be exponentially distributed (Markov) with rate θ . The BSA assigned service time to each RSA is a constant; this is effectively a setup time; a military version is called *target mensuration* and/or *deconfliction*, and refers to the provision of an estimate of the target/threat location by a

human operator using sensor assets plus assurances that friendlies and neutrals are out of range; it should also account for weapon transit time to a target. In medical emergencies it may be initial diagnosis and stabilization of a new patient. We emphasize that in the present case it is immaterial whether a service is known to be completed, successfully or not, before the assigned service completes. Realistically, the BSA has the benefit of follow-up observations, which in military applications is called Battle Damage Assessment (BDA); these are also, realistically, uncertain. Note that if follow-up can be conducted immediately, and service repeated if deemed necessary, the above setup replaces a single service attempt by a quick sequence of assigned services: “Shoot-Look-Shoot”, in military jargon.

2.1 *Fluid or Deterministic Approximation*

Let $R(t)$ be the number of RSAs present in the service region, adopting a simple fluid/deterministic model for the number of RSA tasks in the system at time t ; i.e., $\{R(t), t > 0\}$ is just a real-valued function of continuous time, t . Then, we write

$$\frac{dR(t)}{dt} = \lambda - \theta R(t) - \frac{R(t)}{1 + R(t)} \frac{1}{\tau}, \quad (2.1)$$

where τ is the constant assigned service time, and the effect of service congestion from server saturation is represented by the *Filipiak approximation* (Filipiak, 1988), the term $R(t)/(1 + R(t))$; clearly this term is $\sim R(t)$ for $R(t)$ small, and saturates to ~ 1 as $R(t)$ becomes large, thus reflecting the eventual service rate limitation to $1/\tau$. If desired, the constant parameters in (2.1) can be made time dependent, and time-dependent versions of a problem can be solved numerically, using standard differential equation solvers. There is no convenient closed-form analytical solution for $R(t)$.

Steady-State Success Rate

Let $p(\tau)$ denote the success probability of a task that completes service allocation before loss if τ units of time are allocated to service. Then the success rate satisfies

$$\frac{dS(t)}{dt} = \frac{R(t)}{1+R(t)} \frac{p(\tau)}{\tau}. \quad (2.2)$$

We can use (2.1) to obtain an approximation to the steady-state average number of tasks present in the system by setting the derivative equal to zero and solving for $R = R(\infty)$.

Put $\rho = \lambda \tau$, $\gamma = \theta \tau$,

$$R(\infty) = \frac{2\rho}{1 + \gamma - \rho + \sqrt{(1 + \gamma - \rho)^2 + 4\gamma\rho}}. \quad (2.3)$$

Substituting the above estimate in (2.2) and remembering that tasks arrive at rate λ , the steady-state success probability is estimated by

$$P_F(\lambda, \theta, \tau) = \left(R(\infty) / (1 + R(\infty)) \right) \frac{p(\tau)}{\tau\lambda}. \quad (2.3a)$$

2.2 A Self-Thinning Approximation for an M/G/1 Model of Success Rate

Consider a standard/classical M/G/1 queueing model with a constant service time τ where tasks defect after exponential amount of time with the mean θ^{-1} . Assume defection during service is not observable. It is shown in Gaver *et al* (2000) that a good approximation to the long-run probability that successful service is achieved is

$$P(\lambda, \theta, \tau) = \frac{2}{1 + \rho + \sqrt{(1 + \rho)^2 - 4\rho\delta(\theta)}} p(\tau), \quad (2.4)$$

where $\delta(\theta) = (1 - e^{-\theta\tau}) / \theta\tau$ and $p(\tau)$ is the probability a task given τ time units of service is successfully completed. The traffic intensity is $\rho = \lambda\tau$. An argument for (2.4)

appears in Appendix 2. An analytical/mathematical “exact” approach (forward Kolmogorov equation) is detailed in Appendix 1. The expression (2.4) is remarkably convenient and numerically accurate; however, an improvement is also given in Appendix 2.

Suppose task service times have a distribution F_S but service completion and task loss during service are not observable; then an approximation to the probability of successful task completion for a task that enters service is

$$p(\tau) = \int_0^{\tau} e^{-\theta s} f_S(s) ds \geq e^{-\theta \tau} F_S(\tau) \equiv s(\tau). \quad (2.5)$$

Analytical Model for Single-Service Success Probability

An analytically tractable and flexible expression for $F_S(\tau)$ is the *Fréchet distribution* of extreme value theory (qualitatively appropriate here since it approximates the distribution of maxima); see Resnick (1987):

$$F_S(\tau) = p e^{-(\alpha/\tau)^\beta}, \quad (2.6)$$

where the constant p represents the maximum probability of success, achieved as $\tau \rightarrow \infty$; α is a scale, and β a shape parameter, both positive; if $p < 1$, F_S is a defective distribution. The exponent in the exponential function is unity (1) when $\alpha = \tau$, when the success probability becomes $F_S(\tau) = p e^{-1} = 0.37 p$, or about one-third of the maximum possible, and this independently of β . For $\tau < \alpha$ the exponent increases rapidly as β increases; likewise, it decreases rapidly for $\tau > \alpha$ with increasing p , representing threshold behavior at $\tau = \alpha$.

“Optimum” Single-Server Success Probability

Provided a task enters service, i.e., survives wait in the queue, a lower bound on the probability that it is successful is given by the RHS of (2.5):

$$s(\tau) = e^{-\theta\tau} F_S(\tau) = e^{-\theta\tau} p e^{-(\alpha/\tau)^\beta} \quad (2.7)$$

for the present model of (2.6). There is a unique maximizing value of τ , namely

$$\tau_o = \left(\frac{\beta}{\theta} \right)^{1/(1+\beta)} \alpha^{\beta/(1+\beta)}, \quad (2.8)$$

which implies

$$s^\#(\tau_o) = p \exp \left[-(1+\beta) \left(\frac{\theta\alpha}{\beta} \right)^{\beta/(1+\beta)} \right]. \quad (2.9)$$

Under the policy that all tasks that start service receive τ units of service an approximate lower bound to the probability an arriving task will successfully complete service is given by the right-hand side of

$$P(\lambda, \theta, \tau) \geq \frac{2s(\tau)}{1 + \lambda\tau + \sqrt{(1 + \lambda\tau)^2 - 4\frac{\lambda}{\theta}[1 - e^{-\theta\tau}]}} \equiv \frac{A(\tau; \theta)}{B(\tau; \theta)}. \quad (2.10)$$

The maximizing τ for the lower bound (2.10) with distribution (2.6) can be found numerically. A first order approximation is (2.8). Note that the maximizing τ of (2.8) does not involve λ . An improved approximate maximizing τ can be obtained by applying one iteration of a Newton procedure to $\frac{d}{d\tau} [A(\theta; \tau)/B(\theta; \tau)]$ around the original approximating τ_o . One iteration of Newton's procedure evaluated at τ_o results in

$$\tau_1 = \tau_o - \frac{\frac{d}{d\tau} [A(\theta; \tau_o)/B(\theta; \tau_o)]}{\frac{d^2}{d\tau^2} [A(\theta; \tau_o)/B(\theta; \tau_o)]}. \quad (2.11)$$

The maximizing values of τ from (2.8) for $\alpha = \sqrt{-\ln(0.9)}$ and $\beta = 3$ for various values of θ and λ are displayed in Table 1. Also displayed is the value of τ resulting from one Newton iteration and the resulting lower bounds on the approximate probabilities of an arriving task being successfully served. The table presents the approximate probability of successful service estimated for the previously computed τ by evaluating (2.4). Also displayed are the results from simulation of the queueing system with Fréchet service times and service truncated at the various computed τ . Service completion is not observed and each served task is given τ units of service; service is successful if the service time is less than τ . Each simulation consists of 50 replications; each replication is of 6,000 tasks.

λ	θ	Maximizing τ from (2.8): τ_0 [One iteration of Newton: τ_1]	Lower Bound on the Probability of Success using τ_0 in (2.10) [Lower Bound on the Probability of Success for the Newton τ_1 in (2.10)]	Approximate Probability of Service Success (2.4) for τ_0 [Probability of Service Success; Simulation of system] (Std Error)	Approximate Probability of Service Success, (2.4), for Newton τ_1 [Probability of Service Success; Simulation of System] (Std Error)
0.1	0.25	0.80 [0.78]	0.72 [0.76]	0.85 [0.85] (0.0008)	0.84 [0.84] (0.0007)
	0.50	0.67 [0.66]	0.63 [0.63]	0.74 [0.74] (0.0008)	0.73 [0.74] (0.0009)
0.5	0.25	0.80 [0.69]	0.72 [0.73]	0.80 [0.81] (0.0007)	0.79 [0.79] (0.0007)
	0.50	0.67 [0.61]	0.60 [0.60]	0.70 [0.70] (0.0008)	0.69 [0.68] (0.0009)
1	0.25	0.80 [0.47]	0.64 [0.61]	0.72 [0.74] (0.0009)	0.64 [0.65] (0.001)
	0.50	0.67 [0.50]	0.53 [0.54]	0.62 [0.64] (0.001)	0.58 [0.60] (0.0009)

Table 1

Discussion: The easily computed τ_0 performs reasonably well in all cases computed. The approximate probabilities of successful task completion also agree partly to very well with the simulation results for the actual system using the service policy.

2.3 Towards Discrete Optimal Static Service Responses

Suppose that a BSA has several service options available against the RSA input stream. Service option k is characterized in terms of setup time, τ_k , and corresponding success probability $p_k(\tau_k)$. Under some circumstances it is the practice to take several simultaneous service actions once setup has been performed; under special conditions more than one, e.g., n_k , independent service processes may be applied so as to raise the effective success probability, in which case, replacement of $p_k(\tau_k)$ by $1 - (1 - p_k(\tau_k))^{n_k}$ is justified.

If estimates of task arrival rate, λ , and loss rate, θ , are available, then the optimal static response is to select that option, k^* , that maximizes the probability of surviving without loss to reach the server, surviving service without loss, and actually delivering successful service. Therefore, it is necessary to identify

$$k^* = \text{ArgMax}P(\lambda, \theta, \tau_k)$$

using, say, (2.4), with $p(a)$ replaced with $p_k(\tau_k)$. This then identifies the appropriate service option, which need not be that with highest $p_k(\tau_k)$ value.

Numerical Examples

Here are numerical examples that illustrate the models proposed. Suppose the rate of approach is $\lambda = 0.95$ and the rate of loss is $\theta = 1$. Consider the service options displayed in Table 2.

$k \backslash$	τ_k	p_k	Success Probabilities			
			Fluid	Self-Thinning	Modif. M/G/1	Cond. Prob. Given Start Service
1	0.5	0.8	0.31	0.42	0.43	0.49
2	0.25	0.7	0.43	0.53	0.53	0.55
3	0.15	0.6	0.45	0.51	0.51	0.52

Table 2

Probabilistic model probabilities displayed in Table 2 agree to about two significant digits; they all concur that $k^* = 2$: ($\tau = 0.25$, $p = 0.7$), the intermediate case. This choice remains optimal for the probabilistic models even if task arrival rate drops to 0.75, but the probability of overall success increases.

3. Several ($J \geq 1$) Red Task Types vs. Single Blue Service Agent

Next, consider a single Blue (BSA) confronting a random stream of *different* loss susceptible RSAs. The arrival process of RSAs of type j is Poisson (λ_j) (j is a member of $(1, \dots, J)$) independent of the other task types. Service times for tasks of type j are independent and have a distribution F_j .

Assume the times until loss are independent with those for RSAs of type j having an exponential distribution with mean $1/\theta_j$.

Appendix 3 displays a system of forward Kolmogorov (Takačs-Beneš) equations for the limiting task virtual waiting time for this model. Successive substitution/iteration results in a system of equations for the probability with which a task of type j survives its wait in queue for $j = 1, \dots, J$. Approximate probabilities that an arriving task will start service are also detailed in Appendix 3.

Examples

Although it is possible to obtain the probability an arriving task will start service numerically, the calculation can be tedious for more than two task types. Table 3 displays results of two approximations for the probability an arriving task will start service and the results from a simulation; Bullock (2003). The simulation results are for 50 replications with 6,000 tasks of each type per replication. Both service completion and task loss while in service are observable. Task loss in queue is observable.

Task Type	Task Arrival Rate (Lambda)	Task Loss Rate (Theta)	Means of the Exponential Service Times	Approx. Probability Arriving Task Starts Service (A3-8)	Approx. Probability Arriving Task Starts Service Using Filtered Busy Period (A3-17)	Simulation Fraction of Tasks to Start Service (Std error 50 replications)
1	0.10	1	1	0.78	0.81	0.78 (0.0007)
2	0.10	0.5	3	0.83	0.89	0.84 (0.0007)
3	0.10	0.25	5	0.88	0.97	0.89 (0.0006)
1	0.30	1	1	0.47	0.45	0.46 (0.001)
2	0.30	0.5	3	0.55	0.62	0.57 (0.001)
3	0.30	0.25	5	0.65	0.87	0.70 (0.0009)
1	0.75	1	1	0.22	0.13	0.13 (0.001)
2	0.75	0.5	3	0.28	0.25	0.25 (0.001)
3	0.75	0.25	5	0.37	0.60	0.43 (0.001)

Table 3

Discussion: The first approximation appears to do well for lightly loaded systems. The approximation that thins the arrivals during a busy period appears to do better for heavily

loaded systems. The minimum of the two approximations gives very reasonable agreement with the simulation results for the cases studied.

3.1 Myopic Policies for Choice of Next Task Type to Serve and Amount of Service to Provide

Suppose that loss during service and the completion of service are not observable. A policy is to serve a task of type j for a time τ_j . If the task's service time is less than τ_j , then the task is served successfully; otherwise, it is served unsuccessfully. A successfully served task of type j , results in a reward r_j .

The question: Allocate service time $\tau_j(\bar{n})$ to a type j task as a function of the number of tasks of various types in queue; where $\bar{n} = (n_1, n_2, \dots, n_J)$ with n_j being the number of tasks of type j waiting in queue. Choose the next type of task to serve and the amount of time to give it so as to maximize the long run average reward.

Suppose we allocate $\tau_j > 0$ units of time to a waiting task of type j . The expected reward received is

$$R_E(j, \tau_j; \bar{n}) = r_j \int_0^{\tau_j} e^{-\theta_j y} F_j(dy). \quad (3.1)$$

The expected reward leaving service during the service time is

$$\begin{aligned} R_L(j, \tau_j; \bar{n}) &= r_j + r_j(n_j - 1)A_j(\tau_j; \theta_j) + \sum_{k \neq j} r_k n_k A_j(\tau_j; \theta_k), \\ &+ \sum_{k=1}^J r_k \lambda_k B_j(\tau_j; \theta_k) \end{aligned} \quad (3.2)$$

where

$$A_j(\tau; \theta) = [1 - e^{-\theta \tau}] \quad (3.3)$$

$$B_j(\tau, \theta) = \tau - \frac{1}{\theta} [1 - \exp\{-\theta \tau\}]. \quad (3.4)$$

The myopic policy is to select that j and τ_j , which maximizes the proportion of expected reward stream gained during the next service time.

$$(j_0(\bar{n}), \tau_{j_0}(\bar{n})) = \arg \max_{j, \tau_j} \frac{R_E(j, \tau_j; \bar{n})}{R_L(j, \tau_j; \bar{n})}. \quad (3.5)$$

The policy is myopic because it only optimizes the immediate gain. Similar policies can be obtained for cases in which the service completion is observable and/or the loss of a task during service is observable, etc. Discussion and study of non-myopic policies appear elsewhere.

Examples

All service times and impatience times are independent and exponentially distributed. There are two task types that arrive according to independent Poisson processes. Task losses while waiting and while being served are observed. Task service completion is also observed. In either case, a new task can begin immediately. The myopic policy determines the next task type to be served and the length of the service time to give it. If a task has not completed service when its allocated service time is over, it departs and no reward is collected.

Results from a simulation are shown in Table 4. Each simulation replication is for 100 time units. The number of replications for cases with larger arrival rates is 100. The number of replications for cases with smaller arrival rates is 500. The FIFO policy

serves the first task in the queue until the task is served to completion. The smart FIFO

serves the first task in the queue until time t where τ is the value that maximizes (3.5).

The myopic policy serves that task which maximizes (3.5).

			Reward Task 1=3 Reward Task 2=1			Reward Task 1=1 Reward Task 2=3		
Arrival Rate Task1 [Task2]	Mean Impatience Time Task 1 [Task 2]	Mean Service Time Task 1 [Task 2]	Mean Percent Reward Received FIFO (Std. Error)	Mean Percent Reward Received Smart FIFO (Std. Error)	Mean Percent Reward Received Myopic (Std. Error)	Mean Percent Reward Received FIFO (Std. Error)	Mean Percent Reward Received Smart FIFO (Std. Error)	Mean Percent Reward Received Myopic (Std. Error)
0.1 [0.2]	1 [2]	1 [1.5]	48.48 (0.45)	46.58 (0.46)	48.35 (0.49)	53.04 (0.44)	48.03 (0.44)	52.29 (0.44)
0.25 [0.50]	1 [2]	1 [1.5]	41.96 (0.47)	39.26 (0.45)	43.7 (0.44)	46.5 (0.49)	42.09 (0.5)	46.95 (0.5)
0.5 [1]	1 [2]	1 [1.5]	32.96 (0.35)	32.51 (0.33)	35.18 (0.34)	38.1 (0.34)	34.51 (0.32)	38.1 (0.33)
1 [2]	1 [2]	1 [1.5]	20.8 (0.32)	21.51 (0.3)	25.51 (0.3)	25.33 (0.36)	22.67 (0.3)	25.92 (0.35)
2 [4]	1 [2]	1 [1.5]	8.92 (0.14)	10.46 (0.17)	17.88 (0.21)	13.3 (0.18)	13.69 (0.19)	14.25 (0.18)
4 [8]	1 [2]	1 [1.5]	3.96 (0.06)	4.31 (0.07)	12.54 (0.14)	6.91 (0.09)	7.02 (0.09)	7.08 (0.1)
8 [16]	1 [2]	1 [1.5]	1.87 (0.03)	1.79 (0.02)	7.47 (0.08)	3.57 (0.04)	3.56 (0.04)	3.62 (0.05)
0.1 [0.2]	5 [4]	1 [1.5]	74.21 (0.41)	70.13 (0.41)	73.8 (0.41)	69.19 (0.4)	64.77 (0.43)	69.53 (0.4)
0.25 [0.50]	5 [4]	1 [1.5]	65.33 (0.5)	62.05 (0.48)	66.29 (0.46)	60.27 (0.48)	51.16 (0.47)	61.26 (0.47)
0.5 [1]	5 [4]	1 [1.5]	48.4 (0.4)	47.9 (0.41)	54.35 (0.41)	44.42 (0.41)	42.24 (0.36)	46.27 (0.39)
1 [2]	5 [4]	1 [1.5]	27.86 (0.38)	28.42 (0.38)	41.45 (0.35)	24.09 (0.35)	24.82 (0.32)	28.65 (0.36)
2 [4]	5 [4]	1 [1.5]	14.66 (0.2)	15.11 (0.21)	27.94 (0.25)	11.73 (0.14)	12.44 (0.16)	14.48 (0.18)
4 [8]	5 [4]	1 [1.5]	8.02 (0.09)	8.52 (0.13)	14.71 (0.15)	5.95 (0.07)	6.02 (0.08)	7.1 (0.08)
8 [16]	5 [4]	1 [1.5]	4.1 (0.06)	4.5 (0.06)	7.46 (0.08)	2.79 (0.04)	2.87 (0.04)	3.54 (0.05)

Table 4

Discussion: Not surprisingly, the three policies perform about the same for lightly loaded systems. For lightly loaded systems, FIFO results in larger percentage of rewards because all tasks that start service are given a full service time. For heavier loaded

systems, FIFO and smart FIFO perform about the same. In heavier loaded systems, giving priority to a task type becomes more important.

3.2 A Simulation Study of the Effect of Task Priorities and Service Discipline on Task Completion

In this section, the effect of different task priorities and service disciplines are explored using simulation.

Tasks arrive according to a Poisson process with rate $\lambda = 3.5$. Tasks are of type 1 with probability 0.5 and of type 2 otherwise. A task of type 1 requires service of length $a_1 = 0.2$. A task of type 2 requires a service of length $a_2 = 0.5$. Thus, $\rho = \lambda E[S] = 1.23$ where S is the service time of an arriving task. Each arriving task is lost after an independent random length of time. In this model, losses during service are not observable. Two distributions of task loss times are considered: the uniform and the exponential. All simulation with the same task loss time distributions, use the same simulated arrival times, task types, and task loss times; the difference between the replications is the service discipline and task priority. Table 5 displays the simulated fractions of tasks completed using different service disciplines and task priorities. Each simulation replication consists of 2,000 tasks.

Traffic Intensity=1.23

Distribution of loss times for tasks of type 1	Distribution of loss times for tasks of type 2	Choice of next task to serve	Task type with priority	Frac. of tasks that complete (std. error computed as if observ. are indep.)	Frac. of tasks of type 1 that complete (std. error computed as if observ. are indep.)	Frac. of tasks of type 2 that complete (std. error computed as if observ. are indep.)
Exponential with mean 1	Exponential with mean 3.33	FCFS	2	0.53 (0.01)	0.32 (0.01)	0.73 (0.01)
		FCFS	1	0.64 (0.01)	0.65 (0.02)	0.63 (0.02)
		LCFS	None	0.60 (0.01)	0.50 (0.02)	0.68 (0.01)
		FCFS	None	0.56 (0.01)	0.42 (0.02)	0.69 (0.01)
Uniform on (0.5, 1.5) (2,000 tasks)	Uniform on (2.83, 3.83)	LCFS	1	0.79 (0.01)	0.90 (0.01)	0.68 (0.01)
		FCFS	2	0.58 (0.01)	0.21 (0.01)	0.95 (0.01)
		LCFS	2	0.60 (0.01)	0.30 (0.01)	0.91 (0.01)
		FCFS	1	0.73 (0.01)	0.93 (0.01)	0.51 (0.02)
		LCFS	None	0.71 (0.01)	0.64 (0.02)	0.79 (0.01)
		FCFS	None	0.56 (0.01)	0.23 (0.01)	0.90 (0.01)

Table 5

Discussion: The approximating filtering model for exponential times to loss results in the probability that a task of type 1 survives the queue is equal to 0.626 and the probability a task of type 2 survives the queue is 0.703. For exponential times to loss, the probability a task 1 that starts service is not lost during service is 0.82. The probability a task of type 2 that starts service is not lost during service is 0.86.

Exponential loss times result in a task's position in the queue giving no information on the remaining time until the task is lost. Thus, there is no statistical difference in the fraction of tasks completed successfully between first come first served

(FCFS) and last come first served (LCFS). Since the mean loss time for task 1 is less than that for task 2, task 1 gets priority results in a greater fraction of tasks successfully completed. The myopic policy of Section 4, when the reward for successful completion of both tasks is 1, is to give task 1 priority. Notice that with no priorities, the LCFS discipline outperforms the FCFS discipline in spite of exponentially distributed deadlines. This behavior is a consequence of the different mean impatience times: the longer a task stays in the system, the more likely it is to be of type 2.

Uniform loss times result in a task's position in the queue giving information on the remaining time until the task is lost. The mean loss time for task 1 is less than that for task 2, and task 1 has a shorter service time than task 2. Thus, the best policy is to give task 1 priority with the service discipline LCFS. The phenomena noted for this case also apply for systems with smaller traffic intensities.

4. Concluding Remarks

Modeling uncertain time-critical service systems is a difficult, but vitally important, practical problem. Exact computations are often either impossible or very challenging computationally, especially with multiple task types. Special challenges are present when deciding on a service policy in order to make the system as efficient as possible.

In this paper, we have presented several approximation procedures that are computationally easy and, at least in the examples we have looked at, provide valuable information about the efficiency of the service system under different service options.

An important feature of these approximations is that they stay computationally feasible even for many task types and/or heavily loaded systems.

We have also introduced a heuristic myopic service policy that attempts to maximize *locally* the system efficiency. This policy has performed well under scenarios we have considered.

A number of important issues are left for future work. One such issue is improving the myopic policies into (approximately) optimal policies. A possible approach is a dynamic programming-based procedure that is being developed in Samorodnitsky, Gaver, and Jacobs (2003). Another untouched issue is that of non-stationarity: What happens if the parameters of the system change with time, and need to be constantly estimated in order to update the service policy and keep the system running efficiently? We hope to address the latter question in the near future.

Acknowledgements

The authors wish to acknowledge useful comments by Gideon Weiss, John Hiles, Glen Takahara, and John Lehoczky. John Lehoczky suggested Approximation II in Appendix 2.

References

- Baccelli, F., Boyer, P., and Hebuterne, G., (1984), "Single-server queues with impatient customers," *Adv. Appl. Prob.*, Vol. 16, pp. 887-905.
- Becker, K.J., D.P. Gaver, K.D. Glazebrook, P.A. Jacobs, and S. Lawphongpanich, (2000), "Allocation of tasks to specialized processors: a planning approach," *European Journal of Operational Research*, Vol. 126, pp. 80-88.
- Boots, N.K. and H. Tijms, (1999), "A multi-server queueing system with impatient customers," *Management Science*, **45(3)**, pp. 444-448.
- Boots, N.K. and H. Tijms, (1998), "An M/M/C queue with impatient customers," presented at the First International Workshop on Retrial Queues, Universidad Complutense de Madrid, Madrid, Spain, September 22-24, 1998.
- Brown, M. and F. Proschan, (1983), "Imperfect repair," *J. Appl. Prob.* **20**, pp. 851-859.
- Bullock, G.L. (2003), "UTCT Simulation Model," A simulation implemented in C++.
- Cox, D.R. and W.L. Smith, (1961), *Queues*, Chapman & Hall, London, UK.
- Doytchinov, B., J. Lehoczky, and S. Shreve, (2001), "Real-time queues in heavy traffic with earliest-deadline-first queue discipline," *Ann. Applied Prob.* **11**, pp. 332-378.
- Filipiak, J., (1988), *Modeling and Control of Dynamic Flows in Communication Networks*, Springer-Verlag, Berlin, Germany.
- Gaver, D.P. and P.A. Jacobs, (2002), "BATTLESPACE/INFORMATION WAR (BAT/IW): A System-of-Systems Model of a Strike Operation," Naval Postgraduate School Technical Report, NPS-OR-02-005, Monterey, CA.
- Gaver, D.P. and Jacobs, P.A., (2000), "Servicing impatient tasks that have uncertain outcomes," Naval Postgraduate School Technical Report, NPS-OR-00-001, Monterey, CA.

- Jiang, Z., T.G. Lewis, and J-Y. Colin, (1996), "Scheduling hard real-time constrained periodic tasks on multiple processors," *J. Systems & Software*, **19(11)**, pp. 102-118.
- Kleinrock, L., (1976), *Queueing Systems, Vol. I: Theory*, Wiley (Interscience), New York.
- Lehoczky, J.P., (1996), "Real-time queueing theory," *Proceedings of the IEEE Real-Time Systems Symposium*, December 1996, pp. 186-195.
- Lehoczky, J.P., (1997a), "Using real-time queueing theory to control lateness in real-time systems," *Performance Evaluation Review*, **25(1)**, pp. 158-168.
- Lehoczky, J.P., (1997b), "Real-time queueing network theory," *Proceedings of the IEEE Real-Time Systems Symposium*, December 1997, pp. 58-67.
- Liu, C.L. and J.W. Layland, (1973), "Scheduling algorithms for multiprogramming in a hard real-time environment," *J. Automatic Computing Machinery*, **20(1)**, pp. 40-61.
- Osmundson, J., (2000), "A systems engineering methodology for information systems," *Systems Engineering*, Vol. 3, No. 2, July 2000, pp. 68-81.
- Resnick, S.I., (1987), *Extreme Values, Regular Variation, and Point Processes*, Springer-Verlag, New York.
- Righter, R. (1987) "The stochastic sequential assignment problem with random deadlines," *Probability in the Engineering and the Informational Sciences*, Vol. 1, pp. 189-202.
- Samorodnitsky, G., D.P. Gaver and P.A. Jacobs, (2003), "Choosing the best customer to serve and the ratio control problem," Working paper.
- Ward, A.R. and Bambos, N., "On stability of queueing networks with job deadlines," Technical Report, SU NETLAB-2001-08/01, Engineering Library, Stanford University, Stanford, CA, August 25, 2001.
- Whitt, W., (1999), "Improving service by informing customers about anticipated delays," *Management Science*, **45(2)**, pp. 192-207.

APPENDIX 1

Modified Takaçs-Beneš Equations with Exponential Refusal/Reneging

Let tasks arrive at a service facility according to a Poisson process with rate λ . Service times are independent and identically distributed. Let $W(t)$ be the total virtual work in the system at time t . Each task has a deadline that is exponentially distributed with mean $1/\theta$: if the waiting time or virtual work present when the task arrives exceeds the deadline, the task does not enter the system. This is equivalent to the situation in which tasks whose deadlines have elapsed *when they reach the server* are not served; see e.g., Baccelli *et al* (1984) and Ward *et al* (2001). We will then use a modification to obtain an approximation for the situation in which a deadline may also elapse *during* service.

A.1 Statistically Identified Deadlines and Service to Completion

We start with sketching an argument for derivation of the steady state probability that an arriving task will be successfully served. Let the steady state distribution function of $W(t)$ be

$$F_W(x; t; \theta) = P\{W(t) \leq x\}. \quad (\text{A1-1})$$

A standard renewal theoretical argument shows that $F_W(\cdot; \theta)$ has a right continuous density $p(z; \theta)$ on $(0, \infty)$. Express this as

$$F_W(x; t; \theta) = p_0(t; \theta) + \int_0^x p(z; t; \theta) dz, \quad (\text{A1-1a})$$

where

$$p_0(t; \theta) = P\{W(t) = 0\}. \quad (\text{A1-1b})$$

Since, given $W(t)$, the task joins the queue with probability $e^{-\theta W(t)}$, the probability its deadline does not expire while in queue, one can write

$$F_W(x; \theta) = F_W(x + \Delta t; \theta)[1 - \lambda \Delta t] + \lambda \Delta t \int_0^x (1 - e^{-\theta y} + e^{-\theta y} B(x - y)) F_W(dy; \theta) + o(\Delta t) \quad (\text{A1-2})$$

where B is the distribution function of the positive service time C . Dividing by Δt and letting $\Delta t \downarrow 0$ we obtain

$$p(x; \theta) = \lambda \int_0^x e^{-\theta y} (1 - B(x - y)) F_W(dy; \theta) \quad (\text{A1-3})$$

Use Laplace transforms

$$\psi(s; \theta) = \int_0^\infty e^{-sx} dF_W(x; \theta) \text{ and } b^*(s) = \int_0^\infty e^{-sx} B(dx).$$

Then (A1-3) implies

$$\psi(s; \theta) = p_0(\theta) + \rho \psi(s + \theta; \theta) \delta(s), \quad (\text{A1-4})$$

where

$$\delta(s) = \frac{1 - b^*(s)}{sE[C]}. \quad (\text{A1-5})$$

Substituting $s = 0$ yields $\psi(0; \theta) = 1 = p_0(\theta) + \rho \psi(\theta; \theta)$, hence

$$p_0 = 1 - \rho \psi(\theta; \theta) = 1 - \lambda E[e^{-\theta W}] E[C]. \quad (\text{A1-6})$$

Iterative solution to the equation (A1-4)

Since

$$\psi(s; \theta) = [1 - \rho \psi(\theta; \theta)] + \rho \psi(s + \theta; \theta) \delta(s),$$

it follows that, putting $s = n\theta$ and defining $\psi(\theta) \equiv \psi(\theta; \theta)$, we have

$$\psi(n\theta, \theta) = [1 - \rho \psi(\theta)] + \rho \psi((n+1)\theta; \theta) \delta(n\theta).$$

An inductive argument gives us

$$\begin{aligned}\psi(\theta) = & A(0, \theta) + \dots + A(n, \theta) - \rho\psi(\theta)[A(0, \theta) + \dots + A(n, \theta)] \\ & + \rho\psi((n+1)\theta; \theta)A(n, \theta),\end{aligned}\tag{A1-6a}$$

where

$$\begin{aligned}A(0, \theta) &= 1 \\ A(n, \theta) &= \rho^n \delta(n\theta) \delta((n-1)\theta) \times \dots \times \delta(\theta).\end{aligned}\tag{A1-6b}$$

For $\theta > 0$, $A(n, \theta) \rightarrow 0$. Thus, the probability that an arriving task joins and survives the *queue* before deadline elapse is

$$\psi(\theta) = \frac{\sum_{k=0}^{\infty} A(k, \theta)}{1 + \rho \sum_{k=0}^{\infty} A(k, \theta)}.\tag{A1-7}$$

It is clear that the infinite sum converges faster than exponentially fast for $\theta > 0$, and that this is true for any ρ -value.

Further, for *any* s

$$\psi(s; \theta) = [1 - \rho\psi(\theta)] \sum_{k=0}^{\infty} C(k; s),\tag{A1-8}$$

where

$$C(k; s) = \rho^k \prod_{i=0}^{k-1} \delta(s + i\theta), \quad k \geq 1\tag{A1-9}$$

and

$$C(0; s) = 1.$$

A.2 Services Subject to Detectable Exponential Deadline

If a task deadline's elapse is detectable during service and the task *is then terminated*, then the distribution of service time, \mathbf{C} , must be replaced by that of $\mathbf{C}_T = \min(\mathbf{C}, \text{deadline})$, the *allowed service time*. Consequently, the service times *that*

contribute to the virtual waiting time are, thanks to the exponential deadline assumption, iid with mean

$$E[C_T] = \frac{1 - E[e^{-\theta C}]}{\theta} \quad (\text{A1-10})$$

and tail-transform now

$$\delta_T(s; \theta) = \frac{1 - E[e^{-(\theta+s)C}]}{(\theta+s)E[C_T]} = \frac{\delta(\theta+s)}{\delta(\theta)}. \quad (\text{A1-11})$$

These replace $E[C]$ in ρ , and $\delta(s)$ in the previous solution, (A1-7).

APPENDIX 2

Approximations to Queue Survival for the M/G/1 System with Deadline-Sensitive Delay (The “Self-Thinning” Approximation)

Consider the arrival of tasks with exponentially (mean $1/\theta$) distributed deadlines. Given the virtual waiting time, $W(t)$, at the time of the arrival of a task, the probability that the deadline of the task will not elapse before reaching the server, is simply $e^{-\theta W(t)}$. Instead of letting the task join the queue and then defect if the deadline does elapse before entering service, the same outcome is achieved by simply accepting the task into the queue with probability $e^{-\theta W(t)}$. Based on that, we propose two approximations to the proportion of tasks that are successfully served.

Approximation I

If W has the steady state virtual workload distribution, then the probability that an arriving task is successfully served is $\psi(\theta) = Ee^{-\theta W}$. We neglect the dependence between the fates of different tasks by pretending that the outcomes are decided via a sequence of independent coin tosses with success probability $\psi(\theta)$. The resulting system becomes an M/G/1 queue with traffic intensity $\rho\psi(\theta)$. Then the success probability $\psi(\theta)$ should approximately satisfy the Pollaczek-Khinchine (P-K) formula for M/G/1 queues:

$$\psi(\theta) \equiv E[e^{-\theta W}] = \frac{1 - [\psi(\theta)\lambda]E[C]}{1 - [\psi(\theta)\lambda]E[C] \left\{ \frac{1 - E[e^{-\theta C}]}{\theta E[C]} \right\}}. \quad (\text{A2-1})$$

The simple formula differs somewhat from the solution (A1-7) of the modified Takačs-Beneš equation for the same assumed arrival-queue interaction; but is in handy closed form.

The expression (A2-1) is a quadratic in the desired probability, the solution of which is

$$\psi(\theta) = \frac{2}{1 + \rho + \sqrt{(1 + \rho)^2 - 4\rho\delta(\theta)}}, \quad (\text{A2-2})$$

where $\rho = \lambda E[C]$ as usual, and $\delta(\theta)$ given by (A1-5) is the transform of the service/completion time *tail* or survivor distribution. The approximate probability of successful transit to the server given by this simple expression is unity when $\theta \rightarrow 0$ (no degradation, or infinite deadline), as long as $\rho < 1$; if $\theta \rightarrow \infty$ then, since deadlines are now stringent, the only hope of initiating service is to arrive when there is no server activity, i.e., with probability $1/(1 + \rho)$, and this time any (positive) ρ -value is permitted. In general, there are no restrictions on ρ in (A2-2): a long queue generates many rejections, and thus does not ever remain long, or grow indefinitely. Empirically, the simple expressions, (A2-2) and (A2-9), supply a lower bound that has been shown numerically to be a good approximation to the exact solution of such a reneging or refusal model. Note that the same logic gives as an approximation for the transform of virtual waiting time of non-refused tasks, W , the formula

$$\psi(\xi; \theta) = \frac{1 - \lambda\psi(\theta)E[C]}{1 - \lambda\psi(\theta)E[C] \left\{ \frac{1 - E[e^{-\xi C}]}{\xi E[C]} \right\}}. \quad (\text{A2-3})$$

Approximation II

A refined version of Approximation I accounts for the different experience of a new task that arrives to find the server busy ($W > 0$), as contrasted to one that arrives to find it idle ($W = 0$). Put

$$\psi_+(\theta) = E[e^{-\theta W} | W > 0] \quad (\text{A2-3})$$

the marginal long-run rate of task acceptance given that the server is busy. From the Pollaczek-Khinchine formula

$$\begin{aligned} E[e^{-sW}|W > 0] &= \left[\frac{(1-\rho)\rho\delta(s)}{1-\rho\delta(s)} \right] \frac{1}{\rho} \\ &= \frac{(1-\rho)\delta(s)}{1-\rho\delta(s)}. \end{aligned} \quad (\text{A2-4})$$

We view ρ in this expression as the traffic intensity during a busy period. Then the same logic as the one used in derivation of Approximation I says that the acceptance probability during a busy period $\psi_+(\theta)$ should approximately satisfy the above expression with ρ replaced by $\rho\psi_+(\theta)$. This results in the equation

$$\psi_+(\theta) = (1 - \rho\psi_+(\theta)) \frac{\delta(\theta)}{1 - \rho\psi_+(\theta)\delta(\theta)}. \quad (\text{A2-5})$$

In other words, an auxiliary randomization (biased coin flip) adjusts for the imposition of the deadline, as before in Approximation I, but in a somewhat more refined manner. The solution of (A2-5) is

$$\psi_+(\theta) = \frac{2\delta(\theta)}{(1 + \rho\delta(\theta)) + \sqrt{(1 + \rho\delta(\theta))^2 - 4\rho\delta(\theta)^2}}. \quad (\text{A2-6})$$

For such a ψ_+ -filtered system the expected duration of a busy period, $E[\mathbf{B}]$, should satisfy

$$\begin{aligned} E[\mathbf{B}] &= E[\mathbf{C}] + \rho\psi_+(\theta)E[\mathbf{B}] \\ &= E[\mathbf{C}]/(1 - \rho\psi_+(\theta)). \end{aligned} \quad (\text{A2-7})$$

Consequently, an alternating renewal process argument gives us, as the long-run proportion of time that the server is idle,

$$P\{\mathbf{W} = 0\} = \frac{\lambda^{-1}}{\lambda^{-1} + E[\mathbf{B}]} = \frac{1 - \rho\psi_+(\theta)}{1 + \rho[1 - \psi_+(\theta)]}. \quad (\text{A2-8})$$

Now the probability that an arriving task is admitted (not refused, and eventually served) is

$$\begin{aligned}\psi_2(\theta) &= P\{W = 0\} + (1 - P\{W = 0\})\psi_+(\theta) \\ &= \frac{1}{1 + \rho[1 - \psi_+(\theta)]},\end{aligned}\tag{A2-9}$$

which differs from (A2-2) owing to the more refined conditioning imposed.

Approximation II improves somewhat on Approximation I in all cases explored numerically to date.

APPENDIX 3

Solution to a Modification of the Takaçs-Beneš Equation for Multiple Task Types

Consider a generalization of the model in Appendix 1. Tasks from J task classes arrive to a service facility according to independent Poisson processes with rate λ_j for the j^{th} task class; let $\lambda = \sum_{j=1}^J \lambda_j$. Service times are independent, and service times for each task class are identically distributed. Each task of the j^{th} class has an exponentially distributed deadline with the mean $1/\theta_j$, with the usual independence assumptions. Once again, we start with the case where tasks whose deadlines have elapsed when they reach the server are not served, but no defection occurs while in service. Equivalently, a task whose deadline is shorter than the waiting time at the moment of arrival, does not enter the system. Arguments analogous to those in Appendix 1 show that the Laplace transform of the steady state virtual waiting time in the system satisfies

$$\psi(s; \theta) = p_0(\theta) + \sum_j \rho_j \delta_j(s) \psi(s + \theta_j; \theta), \quad (\text{A3-1})$$

where

$$\delta_j(s) = \frac{1 - b_j^*(s)}{s E[C_j]}, \quad (\text{A3-2})$$

with $\rho_j = \lambda_j E[C_j]$, and $b_j^*(s) = E[e^{-sC_j}]$. Here C_j is a generic service time of a class j task; put $\rho = \sum_{j=1}^J \rho_j$. Furthermore,

$$p_0(\theta) = 1 - \sum_j \rho_j \psi(\theta_j; \theta). \quad (\text{A3-3})$$

An iterative procedure similar to the one used in Appendix 1 shows that the probability $\psi(\theta_j) \equiv \psi(\theta_j; \boldsymbol{\theta})$ that a task of type j will start service (not be lost while in queue) satisfies, for each $n = 1, 2, \dots$, the equation

$$\begin{aligned} \psi(\theta_j) = & A^{(j)}(0, \theta) + \dots + A^{(j)}(n, \theta) \\ & - \left[A^{(j)}(0, \theta) + \dots + A^{(j)}(n, \theta) \right] \sum_{i=1}^J \rho_i \psi(\theta_i) + E_j(n+1; \theta), \end{aligned} \quad (\text{A3-4})$$

where

$$\begin{aligned} A^{(j)}(0, \theta) &= 1 \\ A^{(j)}(n, \theta) &= \sum_{k_1=1}^J \rho_{k_1} \delta_{k_1}(\theta_j) \sum_{k_2=1}^J \rho_{k_2} \delta_{k_2}(\theta_j + \theta_{k_1}) \dots \sum_{k_n=1}^J \rho_{k_n} \delta_{k_n}(\theta_j + \theta_{k_1} + \dots + \theta_{k_{n-1}}) \end{aligned}$$

and

$$E^{(j)}(n; \theta) = \sum_{k_1=1}^J \rho_{k_1} \delta_{k_1}(\theta_j) \dots \sum_{k_n=1}^J \rho_{k_n} \delta_{k_n}(\theta_j + \theta_{k_1} + \dots + \theta_{k_{n-1}}) \psi(\theta_j + \theta_{k_1} + \dots + \theta_{k_{n-1}} + \theta_{k_n})$$

for $n = 1, 2, \dots$. As in Appendix 1, $E^{(j)}(n; \theta) \rightarrow 0$ as $n \rightarrow \infty$, and we obtain a system of linear equations for success probabilities

$$\psi(\theta_j) = \sum_{k=0}^{\infty} A^{(j)}(k; \theta) - \sum_{k=0}^{\infty} A^{(j)}(k; \theta) \sum_{i=1}^J \rho_i \psi(\theta_i) \quad (\text{A3-5})$$

for $j = 1, \dots, J$, which we can solve by replacing the infinite sums by their finite approximations.

A computationally attractive approximation to the solution to the equations (A3-5) is as follows. We start with a self-thinning approximation to the entire aggregation of tasks. Let p be the overall proportion of tasks that start service. Suppose we thin arriving tasks with probability p . Then the transform of the virtual waiting time in the queue is

$$E\left[e^{-sW}\right] = \frac{1 - p \sum_{j=1}^J \lambda_j E[C_j]}{1 - \sum_{j=1}^J \lambda_j p \left[\frac{1 - E\left[e^{-sC_j}\right]}{s} \right]}. \quad (\text{A3-6})$$

On the other hand

$$\begin{aligned} p &= \sum_{j=1}^J \frac{\lambda_j}{\sum_{k=1}^J \lambda_k} E\left[e^{-\theta_j W}\right] \\ &= \sum_{j=1}^J \frac{\lambda_j}{\sum_{k=1}^J \lambda_k} \left[\frac{1 - p\rho}{1 - p \sum_{k=1}^J \rho_k \delta_k(\theta_j)} \right] \equiv B(p). \end{aligned} \quad (\text{A3-7})$$

Note, then $B(p)$ is decreasing in p on $[0, 1/\rho]$ and is always between 0 and 1. Hence, the Equation (A3-7) always has a unique solution \tilde{p} in $[0, 1]$. The approximation for the probability a task of type j starts service is

$$\psi(\theta_j) = E\left[e^{-\theta_j W}\right] = \frac{1 - \tilde{p}\rho}{1 - \tilde{p} \sum_{k=1}^J \rho_k \delta_k(\theta_j)}. \quad (\text{A3-8})$$

A more refined approximation makes use of the fact that the task that arrives when the server is idle always starts service. First of all,

$$E\left[e^{-sW}\right] = P\{W = 0\} + [1 - P\{W = 0\}] E\left[e^{-sW} \mid W > 0\right]. \quad (\text{A3-9})$$

Hence, using the P-K formula for $E[e^{-sW}]$ results in

$$E[e^{-sW} | W > 0] = \frac{1-\rho}{\rho} \frac{\sum_{j=1}^J \rho_j \delta_j(s)}{1 - \sum_{j=1}^J \rho_j \delta_j(s)}. \quad (\text{A3-10})$$

If a proportion p of tasks arriving during a busy period gets to the server, our usual self-thinning approximation results in

$$E[e^{-sW} | W > 0] = \frac{1-\rho p}{1 - \sum_{j=1}^J \rho_j \delta_j(s)} \frac{\sum_{j=1}^J \rho_j \delta_j(s)}{\rho}. \quad (\text{A3-11})$$

Thus, p has to satisfy the relation

$$\begin{aligned} p &= \sum_{j=1}^J \frac{\lambda_j}{\lambda} E[e^{-\theta_j W} | W > 0] \\ &= \sum_{j=1}^J \frac{\lambda_j}{\lambda} \frac{1-\rho p}{1 - \sum_{k=1}^J \rho_k \delta_k(\theta_j)} \frac{\sum_{k=1}^J \rho_k \delta_k(\theta_j)}{\rho}. \end{aligned} \quad (\text{A3-12})$$

The same argument as before shows that this equation has a unique solution \tilde{p}_b in $(0, 1 \wedge \rho^{-1})$.

Let

$$\psi_+(\theta_j) = \sum_{j=1}^J \frac{\lambda_j}{\lambda} \frac{1 - \tilde{p}_b \rho}{1 - \sum_{k=1}^J \rho_k \delta_k(\theta_j)} \frac{\sum_{k=1}^J \rho_k \delta_k(\theta_j)}{\rho} \quad (\text{A3-13})$$

for $j=1, \dots, J$.

The expected length of a busy period satisfies the approximate relation

$$E[\mathbf{B}] = \sum_{j=1}^J \frac{\lambda_j}{\lambda} E[\mathbf{C}_j] + \tilde{p}_b \rho E[\mathbf{B}] \quad (\text{A3-14})$$

and so, approximately,

$$E[\mathbf{B}] = \frac{\sum_{j=1}^J \frac{\lambda_j}{\lambda} E[\mathbf{C}_j]}{1 - \tilde{p}_b \rho}. \quad (\text{A3-15})$$

Since

$$P\{\mathbf{W} = 0\} = \frac{\lambda^{-1}}{\lambda^{-1} + E[\mathbf{B}]} \quad (\text{A3-16})$$

our final approximation is

$$\begin{aligned} \psi_B(\theta_j) &= P\{\mathbf{W} > 0\} \psi_+(\theta_j) + P\{\mathbf{W} = 0\} \\ &= \frac{1 - \tilde{p}_b \rho + \rho \psi_+(\theta_j)}{1 + \rho(1 - \tilde{p}_b)}. \end{aligned} \quad (\text{A3-17})$$

INITIAL DISTRIBUTION LIST

1. Research Office (Code 09).....1
 Naval Postgraduate School
 Monterey, CA 93943-5000

2. Dudley Knox Library (Code 013).....2
 Naval Postgraduate School
 Monterey, CA 93943-5002

3. Defense Technical Information Center2
 8725 John J. Kingman Rd., STE 0944
 Ft. Belvoir, VA 22060-6218

4. Richard Mastowski (Editorial Assistant).....2
 Department of Operations Research
 Naval Postgraduate School
 Monterey, CA 93943-5000

5. Distinguished Professor Donald P. Gaver, Jr.1
 Department of Operations Research
 Naval Postgraduate School
 Monterey, CA 93943-5000

6. Professor Patricia A. Jacobs.....1
 Department of Operations Research
 Naval Postgraduate School
 Monterey, CA 93943-5000

7. Dr. Wendy Martinez1
 Math/Computer/Information Sciences
 Office of Naval Research
 800 North Quincy Street
 Ballston Tower 1, Code 311
 Arlington, VA 22217-5660
 311PS@onr.navy.mil

8. Professor Gennady Samorodnitsky.....1
 School of Operations Research and Industrial Engineering
 Cornell University
 Ithaca, NY 14853

9. Professor John Lehoczky1
 Statistics Department
 Carnegie-Mellon University
 Pittsburgh, PA 15213-3890

10. Professor Kevin Glazebrook1
Management School,
50 George Square
William Robertson Bldg.
Edinburgh University
Edinburgh, UK EH8 9JY
11. Professor Guy Latouche.....1
Universite Libre de Bruxelles
Blvd. d Triomphe Cp. 212
Brussels 1050
Belgium
12. Dr. Guy Fayolle1
Guy.Fayolle@inria.fr